arXiv:1704.04858v1 [stat.ME] 17 Apr 2017

A Nonparametric Bayesian Methodology for Regression Discontinuity Designs

Zach Branson[1], Maxime Rischard[1], Luke Bornn[2], and Luke Miratrix[3]

[1]*Department of Statistics, Harvard University*

[2]*Department of Statistics and Actuarial Science, Simon Fraser University*

[3]*Graduate School of Education and Department of Statistics, Harvard University*

Abstract

Regression discontinuity designs (RDDs) are natural experiments where treatment assignment is determined by a covariate value (or "running variable") being above or below a predetermined threshold. Because the treatment effect will be confounded by the running variable, RDD analyses focus on the local average treatment effect (LATE) *at the threshold*. The most popular methodology for estimating the LATE in an RDD is local linear regression (LLR), which is a weighted linear regression that places larger weight on units closer to the threshold. While LLR exhibits promising bias-reducing properties, LLR tends to yield confidence intervals that undercover, in part because LLR assumes the weighting scheme is fixed, when really there is uncertainty in the weighting scheme choice. We propose an alternative nonparametric methodology utilizing Gaussian process regression that, unlike LLR, (1) does not require specifying a functional form for the expected treatment and control response, and (2) automatically incorporates the uncertainty in how units are weighted when estimating the LATE. We prove our methodology is consistent for the LATE, and we replicate previous simulation studies in the literature to show that our method exhibits better coverage and mean squared error properties than current methodologies.

1 Review of the RDD Problem

Recently there has been a renewed interest in regression discontinuity designs (RDDs), which originated with Thistlethwaite & Campbell (1960). In an RDD, the treatment assignment is discontinuous at a certain covariate value, or "threshold," such that only units whose covariate is above the threshold will receive treatment. There are many examples of RDDs in the applied econometrics literature, such as: the United States providing additional funding to only the 300 poorest counties for the Head Start education program (Ludwig & Miller, 2005); schools mandating students to attend summer school if their exam scores are below a threshold (Matsudaira, 2008); colleges offering financial aid to students whose academic ability is above a cutoff (Van der Klaauw, 2002); and Medicare increasing insurance coverage after age 65 (Card, Dobkin, & Maestas, 2004). The main goal of an RDD is to estimate a treatment effect while addressing likely confounding by the covariate that determines treatment assignment.

We follow the notation of Imbens & Lemieux (2008) and discuss RDDs within the potential outcomes framework: For each unit i, there are two potential outcomes, $Y_i(1)$ and $Y_i(0)$, corresponding to treatment and control, respectively, and background covariate X_i. Only one of these two potential outcomes is observed, but X_i is always observed. Let W_i denote the assignment for unit i, where $W_i = 1$ if i is assigned to treatment and $W_i = 0$ if unit i is assigned to control. The observed outcome for unit i is then

$$y_i = W_i Y_i(1) + (1 - W_i) Y_i(0)$$

Often, one wants to estimate the average treatment effect $\mathbb{E}[Y_i(1) - Y_i(0)]$, but usually this treatment effect is confounded by X (and possibly other unobserved covariates), such that examining the difference in mean response between treatment and control is not appropriate. In these cases, methods such as stratification, matching, and regression are often employed to address covariate confounding when estimating the average treatment effect. However,

such methods are only appropriate when there is sufficient overlap in the treatment and control covariate distributions, i.e.,

$$0 < P(W_i = 1|X_i = x) < 1 \tag{1}$$

For example, this overlap assumption is essential for propensity score methodologies, where the relationship between W_i and X_i is estimated and then accounted for during the analysis.

In an RDD, the relationship between the treatment assignment W_i and the covariates is known. More specifically, it is known that the function $P(W = 1|X = x)$ is discontinuous at some threshold or boundary $X = b$. In this paper we focus on a special case, *sharp RDDs*, where the treatment assignment for a unit i is deterministically

$$W_i = \begin{cases} 1 & \text{if } X_i > b \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

The covariate X that determines treatment assignment in an RDD is often called the "running variable" in order to not confuse it with other background covariates that do not necessarily determine treatment assignment. The more general case, where $P(W = 1|X = x)$ is discontinuous at $X = b$ but is not necessarily 0 or 1, is known as a *fuzzy RDD*.

Rubin (1977) states that when treatment assignment depends on one covariate, estimating $\mathbb{E}[Y_i(1)|X_i]$ and $\mathbb{E}[Y_i(0)|X_i]$ is essential for estimating the average treatment effect. Furthermore, treatment effect estimates are particularly sensitive to model specification of $\mathbb{E}[Y_i(1)|X_i]$ and $\mathbb{E}[Y_i(0)|X_i]$ when there is not substantial covariate overlap, as in a sharp RDD. Aware of this sensitivity, RDD analyses typically do not attempt to estimate the average treatment effect. Instead, they focus on the estimand that requires the least amount of extrapolation to overcome this lack of covariate overlap: the average treatment effect *at*

the boundary b. Defining the conditional expectations for treatment and control as

$$\mu_T(x) \equiv \mathbb{E}[Y_i(1)|X_i = x], \quad \mu_C(x) \equiv \mathbb{E}[Y_i(0)|X_i = x], \tag{3}$$

the treatment effect at the boundary b is (Imbens and Lemieux 2008)

$$\tau = \lim_{x \to b^+} \mathbb{E}[y_i|X_i = x] - \lim_{x \to b^-} \mathbb{E}[y_i|X_i = x] = \mu_T(b) - \mu_C(b). \tag{4}$$

This is often called the local average treatment effect, or LATE (Hahn, Todd, & Van der Klaauw, 2001). The notation $\mu_T(x)$ and $\mu_C(x)$ emphasizes that the goal of an RDD requires estimating two unknown mean functions evaluated at a boundary b. Hahn et al. (2001) (2001) showed that sufficient conditions for τ to be identifiable in a sharp RDD are that $\mu_C(x)$ and $\mu_T(x) - \mu_C(x)$ are continuous at b. They further state that "we can use any nonparametric estimator to estimate" $\mu_T(x)$ and $\mu_C(x)$, and recommended local linear regression (hereafter called LLR), which is a weighted linear regression with weights that are larger for units close to the boundary b.

In Section 2 we review LLR - the most popular nonparametric methodology for estimating the LATE in an RDD - and note its positive and negative properties. While LLR exhibits promising bias-reducing properties, LLR has a well-known tendency to yield confidence intervals that undercover, in part because uncertainty formulae generally assume the weighting scheme in LLR is fixed, when really there is uncertainty in the weighting scheme choice. Furthermore, many practitioners want to estimate the treatment and control response with something more flexible than local linear or quadratic trends. In Section 3, we propose an alternative nonparametric methodology that, unlike LLR, (1) does not require specifying a functional form for the expected treatment and control response, and (2) automatically incorporates the uncertainty in how units are weighted when estimating the treatment effect. In Section 5 we show via simulation that our methodology often exhibits better coverage and RMSE properties than LLR and a variant called robust LLR, which was

developed to improve the coverage of LLR. In Section 6 we conclude by discussing extensions to our methodology to tackle other problems in the RDD literature beyond flexibly fitting $\mu_T(x)$ and $\mu_C(x)$ and improving coverage.

2 Review of Local Linear Regression in RDDs

The goal of an RDD is to well-estimate the LATE τ defined in (4), i.e., to well-estimate $\mu_T(b)$ and $\mu_C(b)$. The LLR methodology in RDDs involves estimating $\mu_T(b)$ as

$$\hat{\mu}_T(b) = X_b(X_T^T W_T X_T)^{-1} X_T^T W_T Y_T \tag{5}$$

where $X_b \equiv \begin{pmatrix} 1 & b \end{pmatrix}$, X_T is the $n_T \times 2$ design matrix corresponding to the intercept and running variable X for treated units, Y_T is the n_T-dimensional column vector of treated units' responses, and W_T is a $n_T \times n_T$ diagonal weight matrix whose entries are

$$(W_T)_{ii} \equiv K\left(\frac{x_i - b}{h}\right), \quad i = 1, \ldots, n_T \tag{6}$$

for some kernel $K(\cdot)$ and bandwidth h. The estimator $\hat{\mu}_C(b)$ is analogously defined for the control. Then, the estimator for the treatment effect is $\hat{\tau} = \hat{\mu}_T(b) - \hat{\mu}_C(b)$. Here, $\hat{\mu}_T(b)$ and $\hat{\mu}_C(b)$ are weighted least squares estimators, where the weights depend on units' closeness to the boundary b. To perform local polynomial regression, one appends higher orders of X to the design matrices X_T and X_C.

2.1 Positives of Local Linear Regression

The RDD literature has largely focused on LLR to estimate the LATE, largely due to LLR's bias-reducing properties. For example, Hahn et al. (2001) recommend LLR over alternatives like kernel regression because Fan (1992) showed that LLR exhibits better bias properties for boundary points than kernel regression. Furthermore, Porter (2003) shows that while kernel

regression exhibits a bias of order $O(h)$ at the boundary, p-order local polynomial regression exhibits a bias of order $O(h^{p+1})$ when the derivatives of $\mu_T(x)$ and $\mu_C(x)$ are sufficiently smooth. For more details on the bias comparison between kernel regression and LLR, please see Imbens & Lemieux 2008 (p. 624-625) as well as Fan & Gijbels (1992).

Finally, LLR's implementation is straightforward once a kernel $K(\cdot)$ and bandwidth h are chosen. The most common choice of $K(\cdot)$ is the rectangular or triangular kernel; Imbens & Lemieux (2008) argue that more complicated kernels rarely make a difference in estimation. Much more attention has been given to the bandwidth choice h, largely because the bias is characterized by h. In the 2000s, choosing an appropriate h for LLR in RDDs was an open problem: For example, Ludwig & Miller (2005) stated that "there is currently no widely-agreed-upon method for selection of optimal bandwidths...so our strategy is to present results for a broad range of candidate bandwidths." However, the current state-of-the-art bandwidth selection method is that of Imbens & Kalyanaraman (2011), who derived a data-driven, MSE-optimal bandwidth for LLR estimators. This provided practitioners with clear guidelines for implementing LLR for RDDs, which made its use very popular.

With these promising results in mind, there are nonetheless two notable concerns with the LLR methodology for RDDs, and we discuss them below. These concerns have been discussed previously, but they have not been fully addressed by current RDD methodologies. Addressing these two concerns is the motivation for our methodology outlined in Section 3.

2.2 Concern One: LLR Does Not Allow a Flexible Fit to $\mu_T(x)$ and $\mu_C(x)$

While much of the RDD literature has focused on first- or sometimes second-order local polynomial methodologies, many practitioners have used high-order global polynomials to more flexibly fit the mean response functions $\mu_T(x)$ and $\mu_C(x)$. One can see the varied use of polynomial regression methods by referring back to the applied RDD papers mentioned at the beginning of this paper: Matsudaira (2008) argued for a global third-order polynomial

regression, and also considered fourth- and fifth-order polynomials as a sensitivity check; similarly, Van der Klaauw (2002) used a global third-order polynomial and noted that LLR could have been an alternative; finally, Card et al. (2004) argued for using a global third-order polynomial regression instead of LLR because the running variable, age, was discrete.

On the other hand, many have argued against the use of high-order polynomials in RDDs. Possibly the most notable of these arguments,"Why High-Order Polynomials Should Not Be Used In Regression Discontinuity Designs" (Gelman & Imbens, 2014), demonstrates that point estimates and confidence intervals estimated from high-order polynomials are highly sensitive to the order of the polynomial, and thus recommends local linear or quadratic regression instead. Nonetheless, there is a desire in the literature for models that flexibly estimate the underlying functions $\mu_T(x)$ and $\mu_C(x)$ but also yield trustworthy inferences.

2.3 Concern Two: LLR Does Not Incorporate Uncertainty in the Bandwidth Choice

The bandwidth is arguably the most important choice to be made in the LLR methodology for RDDs, because the treatment effect is often sensitive to the bandwidth choice. This motivates sensitivity checks such as that in Ludwig & Miller (2005), where the treatment effect is estimated several times with different bandwidths to ensure that estimates do not vary too greatly. This process showcases that there is often uncertainty in the bandwidth choice. However, uncertainty estimates for the treatment effect in LLR methodologies are characterized by plugging in a single bandwidth, whether it is the MSE-optimal bandwidth of Imbens & Kalyanaraman (2011) or a bandwidth selected by another methodology such as cross-validation. Current LLR methodologies do not attempt to capture uncertainty in the bandwidth choice.

This may be one reason why LLR tends to exhibit undercoverage. Calonico et al. (2014) attempted to fix the undercoverage issue with a methodology hereafter called "robust LLR" that shifts point estimates and inflates confidence intervals from LLR by incorporating a bias

correction estimated by a "robust bandwidth" that is larger than the typical bandwidths used for estimation. Calonico et al. (2014) showed via simulation that robust LLR exhibits better coverage than LLR, and thus it is quickly becoming the state-of-the-art for inference in RDDs. However, robust LLR still does not attempt to capture the uncertainty in the bandwidth choice, which may explain why it still tends to undercover in simulation.

2.4 Addressing These Concerns with Gaussian Process Regression

In Section 3 we propose using one of the most common methodologies for flexibly estimating unknown functions $\mu_T(x)$ and $\mu_C(x)$: Gaussian Process Regression (GPR). Gaussian processes are commonly used as a general prior on unknown functions (Rasmussen & Williams, 2006), and GPR is a nonparametric model that does not require specifying a functional form for $\mu_T(x)$ and $\mu_C(x)$. First, we review how Gaussian processes are used to estimate a single unknown function, and thus why they are a natural option for RDDs, which require estimating two unknown functions. To the best of our knowledge, Gaussian processes have not previously been proposed for RDDs, and the notion of using GPR to model $\mu_T(x)$ and $\mu_C(x)$ generates a large class of possible models. In Section 3 we outline a roadmap of assumptions commonly placed on $\mu_T(x)$ and $\mu_C(x)$ which in turn suggest different GPR models. For each GPR model we provide the posterior of the LATE and compare it to the analogous LLR model. Furthermore, we outline an extension to our GPR models that captures the uncertainty in how units are weighted when estimating the LATE. In Section 4 we show that the posteriors from our GPR models are consistent in estimating the treatment effect. In Section 5 we show that our GPR method often exhibits better coverage and RMSE properties than both LLR and robust LLR.

3 GPR Models to Estimate the LATE τ

First we review notation for GPR and how GPR is used to estimate a single unknown function. We then discuss several GPR models that estimate the two unknown mean functions in sharp RDDs.

3.1 Notation for Estimating One Unknown Function

Define a dataset $\{x_i, y_i\}_{i=1}^n$ of responses $\mathbf{y} = (y_1, \ldots, y_n)$ that varies around some unknown function of the covariate $\mathbf{x} = (x_1, \ldots, x_n)$:

$$\mathbf{y} = f(\mathbf{x}) + \epsilon \tag{7}$$

where $\epsilon \sim \mathcal{N}_n(0, \sigma_y^2 \mathbf{I}_n)$. If the goal is to well-estimate $\mathbb{E}[f(x^*)]$ for a particular covariate value x^*, one option is to specify a functional form for $f(\mathbf{x})$ and then predict $\mathbb{E}[f(x^*)]$ from this specified model, such as local linear regression, as discussed in Section 2. Instead of specifying a functional form for $f(\mathbf{x})$, we consider nonparametrically inferring $f(\mathbf{x})$ by placing a prior on $f(\mathbf{x})$. A Gaussian process is one such prior:

$$f(\mathbf{x}) \sim \text{GP}(m(\mathbf{x}), K(\mathbf{x}, \mathbf{x}')) \tag{8}$$

for some unknown mean function $m(\mathbf{x})$ and covariance function $K(\mathbf{x}, \mathbf{x}')$. The notation $f(\mathbf{x}) \sim GP(m(\mathbf{x}), K(\mathbf{x}, \mathbf{x}'))$ denotes a Gaussian process prior on the unknown function $f(\mathbf{x})$, which states that, for any $(x_1, \ldots, x_n)$, the joint distribution $(f(x_1), \ldots, f(x_n))$ is an n-dimensional multivariate normal distribution with mean vector $(m(x_1), \ldots, m(x_n))$ and covariance matrix $K(\mathbf{x}, \mathbf{x}')$ whose (i, j) entries are $K(x_i, x_j')$.

There are many choices one could make for the mean and covariance functions. A common choice for the mean function is $m(\mathbf{x}) = 0$; a common choice for the covariance function is the

squared exponential covariance function, whose entries are

$$K(x_i, x_j) \equiv \sigma^2_{GP} \exp\left(-\frac{1}{2\ell^2}(x_i - x_j)^2\right). \tag{9}$$

Placing a Gaussian process prior with a squared exponential covariance function on $f(\mathbf{x})$ assumes that $f(\mathbf{x})$ is infinitely differentiable, which is similar to other assumptions in the RDD literature (e.g., Assumption 3.3 of Imbens & Kalyanaraman 2011 and Assumption 1 of Calonico et al. 2014). The covariance parameters σ^2_{GP} and ℓ are called the variance and lengthscale, respectively. The variance determines how much an estimate of $f(\mathbf{x})$ can vary from its mean. The lengthscale determines the smoothness of the function: Small lengthscales correspond to $f(\mathbf{x})$ changing rapidly. Most importantly, the covariance function assumes that the response at a particular covariate value $f(x^*)$ will be similar to the response at covariate values close to x^*. This makes the covariance parameters similar to the kernel and bandwidth in the LLR methodology.

In practice, the covariance parameters are estimated from the data, such as through maximum likelihood or cross-validation Rasmussen & Williams 2006 (Chapter 5). In Section 3.2 we assume that the covariance parameters are fixed, and in Section 3.3 we extend to a full Bayesian approach that incorporates uncertainty in these parameters.

3.2 Three GPR Models for Sharp RDDs

The notion of using GPR to estimate the LATE in a sharp RDD suggests a class of models that—to the best of our knowledge—has not previously been considered in the RDD literature. Here we present three different GPR models, depending on the assumptions that are placed on the unknown response functions $\mu_T(x)$ and $\mu_C(x)$. For each model, we show the resulting posterior for the LATE and compare it to its analogous LLR model. In Section 3.3 we discuss how—unlike LLR methodologies—the uncertainty in how units are weighted can be incorporated into these three GPR models.

For all three GPR models, we assume that the treatment response $Y_i(1)$ and the control response $Y_i(0)$ have the following relationship with the running variable x_i

$$\begin{aligned} &Y_i(1) = \mu_T(x_i) + \epsilon_{i1}, \quad \text{and} \quad Y_i(0) = \mu_C(x_i) + \epsilon_{i0}, \quad \text{where} \\ &\epsilon_{i1} \sim N(0, \sigma_{y1}^2) \quad \text{and} \quad \epsilon_{i0} \sim N(0, \sigma_{y0}^2) \end{aligned} \tag{10}$$

The goal of an RDD is to best estimate the two unknown mean functions $\mu_T(x)$ and $\mu_C(x)$ evaluated at the boundary b. Thus, intuitively, the procedure outlined in Section 3.1 can simply be performed twice. However, there are assumptions on $\mu_T(x)$ and $\mu_C(x)$ that, if true, can simplify our GPR models and make inference more precise. Below we present three assumptions in order of restrictiveness; each assumption corresponds to a different GPR model. In all three models, we assume that $\mu_T(x)$ and $\mu_C(x)$ are stationary processes, i.e., the covariance parameters of $\mu_T(x)$ and $\mu_C(x)$ do not vary with the running variable X. We discuss cases when this stationarity assumption is inappropriate in Sections 5 and 6.

Mean Shift Assumption: $\mu_T(x) = \mu_C(x) + \tau$ for some constant τ, for all $x \in \mathbb{R}$, and $\mu_T(x)$ and $\mu_C(x)$ are stationary processes.

The Mean Shift Assumption states that the mean treatment and control response functions are identical up to some constant shift τ. In this case, the proper LLR methodology would be to fit a local linear regression on both sides of the boundary with the same bandwidth and slope but different intercepts; then, the difference in the estimated intercepts would be the treatment effect estimate. Instead of specifying a functional form for $\mu_C(x)$ (and consequentially $\mu_T(x)$ by assumption), we place a Gaussian process prior on $\mu_C(x)$:

$$\mu_C(x) \sim \text{GP}(m(\mathbf{x}), K(\mathbf{x}, \mathbf{x}'))$$

and a prior for τ of

$$\tau \sim N(0, \sigma_\tau^2)$$

for some large σ_τ^2, unless there is prior information on the treatment effect.

Given the above priors, the posterior for τ is

$$\tau|\mathbf{x}, \mathbf{y} \sim N\left(\text{Cov}(\mathbf{y}, \tau)^T \text{Var}(y)^{-1}(\mathbf{y} - m(\mathbf{x})), \sigma_\tau^2 - \text{Cov}(\mathbf{y}, \tau)^T \text{Var}(y)^{-1} \text{Cov}(\mathbf{y}, \tau)\right) \tag{11}$$

where

$$\begin{aligned} \text{Var}(y) &= K(\mathbf{x}, \mathbf{x}') + \sigma_\tau^2 \mathbb{I}(\mathbf{x} > b)\mathbb{I}(\mathbf{x} > b)^T + \sigma_y^2 \mathbf{I}_n \\ \text{Cov}(\mathbf{y}, \tau) &= \sigma_\tau^2 \mathbb{I}(\mathbf{x} > b) \end{aligned}$$

where $\mathbb{I}(A)$ is the indicator function for event A, such that $\mathbb{I}(\mathbf{x} > b)$ is a n-dimensional column vector with n_C zeroes and n_T ones. The derivation of the posterior (11) is in the Appendix. We call this resulting posterior the "shifted GP model."

If the Mean Shift Assumption holds, fitting local linear regressions on both sides of the boundary with the same bandwidth and slope but different intercepts is beneficial in the sense that more data will be used to estimate the bandwidth and slope, yielding more precise estimates. However, if the Mean Shift Assumption does not hold, then this procedure will likely yield undesireably biased estimates. The same cautionary tale holds for the shifted GP model: More data would be used to estimate the covariance parameters of the underlying mean function $\mu_C(x)$, but this can lead to undesireable results if the Mean Shift Assumption does not hold.

Note that the Mean Shift Assumption—though a strong assumption—is indistinguishable from the assumption that all of the derivatives of $\mu_T(x)$ and $\mu_C(x)$ at the boundary $x = b$ are equal, because $\mu_T(x)$ is only observed for $x > b$ and $\mu_C(x)$ for $x \leq b$. Such an assumption

on the derivatives of the mean treatment and control response functions is similar to other assumptions found in the RDD literature (e.g., Assumption 3.3 of Imbens & Kalyanaraman 2011).

Instead of placing restrictions on both the mean and covariance functions of treatment and control, one can instead place a restriction on only the covariance functions.

Same Covariance Assumption: $\text{Cov}(\mu_T(x)) = \text{Cov}(\mu_C(x))$, and $\mu_T(x)$ and $\mu_C(x)$ are stationary processes.

If the Same Covariance Assumption holds, a natural LLR procedure is to fit local linear regressions on both sides of the boundary with the same bandwidth but different intercepts and slopes. This is largely the standard practice in the RDD literature (Imbens & Lemieux, 2008).

Instead of specifying a functional form for the unknown functions $\mu_T(x)$ and $\mu_C(x)$, we place Gaussian process priors on $\mu_T(x)$ and $\mu_C(x)$ for given mean functions $m_T(\mathbf{x})$ and $m_C(\mathbf{x})$ and covariance function $K(\mathbf{x}, \mathbf{x}')$:

$$\begin{aligned} \mu_T(x) &\sim \text{GP}(m_T(\mathbf{x}), K(\mathbf{x}, \mathbf{x}')) \\ \mu_C(x) &\sim \text{GP}(m_C(\mathbf{x}), K(\mathbf{x}, \mathbf{x}')) \end{aligned} \tag{12}$$

Unlike the model under the Mean Shift Assumption, this model does not explicitly have a parameter corresponding to the treatment effect. Instead, estimates $\hat{\mu}_T(b)$ and $\hat{\mu}_C(b)$ are obtained, which results in a treatment effect estimate $\hat{\tau} = \hat{\mu}_T(b) - \hat{\mu}_C(b)$, analogous to the LLR methodology.

Given the Gaussian process prior mean functions $m_T(\mathbf{x})$ and $m_C(\mathbf{x})$ and covariance function $K(\mathbf{x}, \mathbf{x}')$, as well as their parameters, the following posteriors for $\mu_T(b)$ and $\mu_C(b)$ can be obtained via standard conditional multivariate normal theory (for an exposition,

please see Rasmussen & Williams 2006, Pages 16-17):

$$\begin{aligned} \mu_T(b)|\mathbf{x},\mathbf{y} &\sim N(\mu_{T|b}, \sigma^2_{GP} - \Sigma_{T|b}) \\ \mu_C(b)|\mathbf{x},\mathbf{y} &\sim N(\mu_{C|b}, \sigma^2_{GP} - \Sigma_{C|b}), \quad \text{where} \end{aligned} \tag{13}$$

$$\begin{aligned} \mu_{T|b} &\equiv m_T(b) + K(b,\mathbf{x}_T)[K(\mathbf{x}_T,\mathbf{x}_T) + \sigma_y^2\mathbf{I}]^{-1}(\mathbf{y}_T - m_T(\mathbf{x}_T)) \\ \mu_{C|b} &\equiv m_C(b) + K(b,\mathbf{x}_C)[K(\mathbf{x}_C,\mathbf{x}_C) + \sigma_y^2\mathbf{I}]^{-1}(\mathbf{y}_C - m_C(\mathbf{x}_C)) \\ \Sigma_{T|b} &\equiv K(b,\mathbf{x}_T)[K(\mathbf{x}_T,\mathbf{x}_T) + \sigma_y^2\mathbf{I}]^{-1}K(\mathbf{x}_T,b) \\ \Sigma_{C|b} &\equiv K(b,\mathbf{x}_C)[K(\mathbf{x}_C,\mathbf{x}_C) + \sigma_y^2\mathbf{I}]^{-1}K(\mathbf{x}_C,b) \end{aligned} \tag{14}$$

Here, $\mu_{T|b}$ and $\mu_{C|b}$ denote the posterior mean for $\mu_T(b)$ and $\mu_C(b)$, respectively, which are the estimands in the definition of the local treatment effect τ, as defined in (4). Note that $\mu_{T|b}$ and $\mu_{C|b}$ are weighted averages of the observed response, where the weights $K(b,\mathbf{x})[K(\mathbf{x},\mathbf{x}) + \sigma_y^2\mathbf{I}]^{-1}$ depend on the covariance parameters σ^2_{GP}, ℓ, and σ^2_y. In Section 3.3 we extend our model to incorporate uncertainty in these parameters and thus uncertainty in how units are weighted when estimating the LATE. For more discussion on the behavior of the weights $K(b,\mathbf{x})[K(\mathbf{x},\mathbf{x}) + \sigma_y^2\mathbf{I}]^{-1}$, please see Rasmussen & Williams 2006 (Section 2.6).

The posterior for the treatment effect under the Same Covariance Assumption is then

$$\tau \equiv \mu_T(b) - \mu_C(b)|\mathbf{x},\mathbf{y} \sim N\left(\mu_{T|b} - \mu_{C|b}, 2\sigma^2_{GP} - \Sigma_{T|b} - \Sigma_{C|b}\right). \tag{15}$$

We call this resulting posterior the "two-GP model" with the same covariance in treatment and control.

Finally, if neither the Mean Shift or Same Covariance Assumptions hold, one can still assume that the mean treatment and control response processes are stationary, but allow both the mean and covariance to vary on either side of the boundary.

Stationary Assumption: $\mu_T(x)$ and $\mu_C(x)$ are stationary processes.

The posterior in this case would be identical to (15), except the means $\mu_{T|b}$ and $\mu_{C|b}$ and covariances $\Sigma_{T|b}$ and $\Sigma_{C|b}$ are instead defined as

$$\begin{aligned}
\mu_{T|b} &\equiv m_T(b) + K_T(b, \mathbf{x}_T)[K_T(\mathbf{x}_T, \mathbf{x}_T) + \sigma_{y1}^2 \mathbf{I}]^{-1}(\mathbf{y}_T - m_T(\mathbf{x}_T)) \\
\mu_{C|b} &\equiv m_C(b) + K_C(b, \mathbf{x}_C)[K_C(\mathbf{x}_C, \mathbf{x}_C) + \sigma_{y0}^2 \mathbf{I}]^{-1}(\mathbf{y}_C - m_C(\mathbf{x}_C)), \\
\Sigma_{T|b} &\equiv K_T(b, \mathbf{x}_T)[K_T(\mathbf{x}_T, \mathbf{x}_T) + \sigma_{y1}^2 \mathbf{I}]^{-1} K_T(\mathbf{x}_T, b) \\
\Sigma_{C|b} &\equiv K_C(b, \mathbf{x}_C)[K_C(\mathbf{x}_C, \mathbf{x}_C) + \sigma_{y0}^2 \mathbf{I}]^{-1} K_C(\mathbf{x}_C, b)
\end{aligned} \tag{16}$$

i.e., the shared covariance $K(\cdot, \cdot)$ is replaced with $K_T(\cdot, \cdot)$ for units receiving treatment and $K_C(\cdot, \cdot)$ for units receiving control. The analogous LLR methodology would be to allow different intercepts, slopes, and bandwidths on either side of the boundary, but this is rarely done in practice. For example, Imbens & Lemieux (2008) argue that if the curvature of $\mu_T(x)$ and $\mu_C(x)$ is the same, then the large-sample optimal bandwidths should be the same; and, furthermore, there is additional noise in estimating two optimal bandwidths rather than one, due to the smaller sample used to estimate each bandwidth. A similar argument holds for the two-GP model: The notion of $\mu_T(x)$ and $\mu_C(x)$ having similar curvature is analogous to the notion of $\mu_T(x)$ and $\mu_C(x)$ having similar covariance functions.

In summary, we have presented three Gaussian process models to estimate the LATE in a sharp RDD:

1. The shifted GP model
2. The two GP model with the same covariance in treatment and control
3. The two GP model with different covariance in treatment and control

Their associated posteriors assume fixed mean and covariance parameters, and in practice maximum-likelihood or cross-validation could be used for estimation; this is analogous to plugging in a single estimate for the slope and bandwidth in LLR. We next show how this

can be extended to incorporate uncertainty in the mean and covariance parameters; this is analogous to incorporating uncertainty in both the slope and bandwidth in LLR. Typical uncertainty estimates for the LATE in the LLR methodology account for uncertainty in the slope but not the bandwidth.

3.3 Accounting for Mean and Covariance Function Uncertainty

All three GPR models presented in Section 3.2 assume that $\mu_T(b)$ and $\mu_C(b)$ will be similar to $\mu_T(x)$ and $\mu_C(x)$, respectively, for x near b. The extent of this similarity is determined by the covariance parameters of $K(\mathbf{x}, \mathbf{x}')$; e.g., if $K(\mathbf{x}, \mathbf{x}')$ is the squared exponential covariance function defined in (9), then this similarity is determined by σ^2_{GP}, and ℓ. This is analogous to the weights determined in LLR: If the response at the boundary is not very similar to responses relatively near the boundary, then the weights should decay rapidly for units further from b. Thus, incorporating uncertainty in the covariance parameters is analogous to incorporating uncertainty in the kernel and bandwidth in LLR.

We now extend these models to a full Bayesian approach such that uncertainty in these parameters is incorporated when estimating the LATE. Denote the mean function parameters by $\boldsymbol{\theta}_m$ and the covariance function parameters by $\boldsymbol{\theta}_K$. For example, consider the two-GP model with mean functions

$$\begin{aligned} m_T(\mathbf{x}) &= \mathbf{h}(x)^T \boldsymbol{\beta}_T \\ m_C(\mathbf{x}) &= \mathbf{h}(x)^T \boldsymbol{\beta}_C \end{aligned} \tag{17}$$

where $\mathbf{h}(x) = (1, x, \ldots, x^{p-1})$, and $\boldsymbol{\beta}_T$ and $\boldsymbol{\beta}_C$ are the corresponding p-dimensional column vectors. In this case, $\boldsymbol{\theta}_m = (\boldsymbol{\beta}_T, \boldsymbol{\beta}_C)$. For the squared exponential covariance function defined in (9), $\boldsymbol{\theta}_K = (\sigma^2_y, \sigma^2_{GP}, \ell)$.

In order to incorporate uncertainty in $\boldsymbol{\theta_m}$ and $\boldsymbol{\theta_K}$, one can first obtain draws $1, \ldots, D$ from the joint posterior of $(\boldsymbol{\theta}_m, \boldsymbol{\theta}_K)$, rather than obtaining point-estimates $\hat{\boldsymbol{\theta}}_m$ and $\hat{\boldsymbol{\theta}}_K$. Then, for each draw $(\boldsymbol{\theta}_m, \boldsymbol{\theta}_K)_1, \ldots, (\boldsymbol{\theta}_m, \boldsymbol{\theta}_K)_D$, one draws from the posterior (11) or the posterior (15), depending on which assumptions in Section 3.2 hold.

Section 3.2 already defines the likelihood for the shifted GP and two-GP models, so all that remains is to specify priors for $(\boldsymbol{\theta}_m, \boldsymbol{\theta}_K)$ in order to obtain draws from the joint posterior of $(\boldsymbol{\theta}_m, \boldsymbol{\theta}_K)$. Priors for $(\boldsymbol{\theta}_m, \boldsymbol{\theta}_K)$ will depend on the choice of mean and covariance functions. For example, for the mean functions defined in (12), we recommend $\mathcal{N}(\mathbf{0}, B)$ as a prior for $\boldsymbol{\beta}_T$ and $\boldsymbol{\beta}_C$, where B is a $p \times p$ diagonal matrix with reasonably large entries (Rasmussen & Williams 2006, Pages 28-29). For the squared exponential covariance function given by (9). We recommend half-Cauchy priors for the covariance parameters $\frac{1}{\ell^2}$ and σ^2_{GP} and noise σ^2_y, following advice from Gelman et al. (2006) about stable priors for variance parameters.

Although one could plug in estimates for the mean and covariance parameters, we recommend this full Bayesian approach because it incorporates uncertainty in how units are weighted when estimating the LATE in addition to providing a flexible fit for the mean response functions $\mu_T(x)$ and $\mu_C(x)$. In Section 5, we show via simulation that this full Bayesian approach tends to exhibit better coverage and RMSE properties than LLR and the robust LLR methodology proposed by Calonico et al. (2014).

4 Posterior Consistency of our GPR Models

Gaussian processes are known to exhibit posterior consistency under minimal assumptions. Ghosal & Roy (2006) proved posterior consistency of binary GPR for fixed covariance parameters, and Choi & Schervish (2007) proved posterior consistency of GPR when the response is continuous. More generally, van der Vaart & van Zanten (2008) studied the contraction rate for Gaussian process priors for density estimation and regression problems, and van der Vaart & van Zanten (2009) extended these results to when a prior is placed on the length-scale of a Gaussian process. All three of our models in Section 3.2 estimate the treatment effect as the difference between two Gaussian process regressions; thus, our posterior of the treatment effect is consistent if the separate GPRs on either side of the discontinuity are consistent. First we show posterior consistency assuming the covariance parameters are fixed, as

in Section 3.2, and then we extend these results to when a prior is placed on the covariance parameters, as in Section 3.3.

We prove posterior consistency assuming the Stationary Assumption in Section 3.2, but the results also hold for the Same Covariance Assumption and Mean Shift Assumption. Furthermore, we assume the mean functions $m_T(x) = 0$, $m_C(x) = 0$ and the squared exponential covariance function defined in (9). Discussion about the extent to which these results extend to other choices for the mean and covariance functions is in the Appendix. The other assumptions necessary for Theorems 1 and 2 below follow van der Vaart & van Zanten (2009) and are also given in the Appendix.

Theorem 1: *Assume that the Stationary Assumption holds, the covariance functions $K_T(x, x)$ and $K_C(x, x)$ are fixed, and Assumptions A1, A2, and A3 given in the Appendix hold. Denote the true LATE as $\tau^* = \mu_T^*(b) - \mu_C^*(b)$, where $\mu_T^*(x)$ and $\mu_C^*(x)$ are the true mean treatment and control response functions in the model (10). Let $\prod(\tau|x_1, \ldots, x_n)$ denote the posterior distribution of the LATE, defined in (15). Then, this posterior is consistent, in the sense that*

$$\prod (\tau : h(\tau, \tau^*) \geq M\epsilon_n | x_1, \ldots, x_n) \xrightarrow{P_{\tau^*}} 0 \tag{18}$$

for sufficiently large M, where h is the Hellinger distance, and ϵ_n is the rate at which the posterior of τ contracts to the true τ^*.

The proof of Theorem 1, as well as the nature of the contraction rate, is given in the Appendix. Theorem 2 establishes posterior consistency when a prior is placed on the lengthscale parameter ℓ, instead of being held fixed (see Section 3.3). A discussion about posterior consistency when an additional prior is placed on σ_{GP}^2 is in the Appendix.

Theorem 2: *Assume that the Stationary Assumption holds, the σ^2_{GP} parameters in $K_T(x,x)$ and $K_C(x,x)$ are fixed, and Assumptions A1, A2, A3, and A4 given in the Appendix hold. Then, Theorem 1 holds.*

The proof of Theorem 2 is given in the Appendix. A corollary follows from the proofs of Theorem 1 and 2.

Corollary 1: *Theorems 1 and 2 hold if the Same Covariance Assumption or Mean Shift Assumption hold instead of the Stationary Assumption.*

5 Simulation Results

Here we investigate how our Gaussian Process model compares to LLR and the robust LLR method introduced in Calonico et al. (2014). We choose these two methods because the former is the standard in both applied work and the RDD literature at large, and the latter is the most recent method that attempts to solve the undercoverage issue of LLR. We focus on the two-GP model assuming the Same Covariance Assumption in Section 3.2 and use the mean functions

$$\begin{aligned} m_T(\mathbf{x}) &= \boldsymbol{\beta}_{0T} + \boldsymbol{\beta}_{1T}\mathbf{x} \\ m_C(\mathbf{x}) &= \boldsymbol{\beta}_{0C} + \boldsymbol{\beta}_{1C}\mathbf{x} \end{aligned} \tag{19}$$

and the squared exponential covariance function given by (9). These assumptions are analogous to the LLR procedure of fitting separate local linear regressions in treatment and control with differing slopes but the same bandwidth. We placed half-Cauchy priors on the covariance parameters as discussed in Section 3.3, and we used the `R` package `rstan` to sample from the posterior of these covariance parameters.

We conduct a simulation study based on those found in the RDD literature. In all simulations, we generated 1,000 datasets of 500 observations $\{(x_i, y_i, \epsilon_i) : i = 1, \ldots, 500\}$,

where $x_i \sim 2\text{Beta}(2,4) - 1$, $\epsilon_i \sim N(0, 0.1295^2)$, and $y_i = \mu_j(x_i) + \epsilon_i$ for different mean functions $\mu_j(x_i)$. This setup follows previous simulations seen in Imbens & Kalyanaraman (2011) and Calonico et al. (2014).

We consider seven different mean functions (see Figure 1), which we call Lee, Quad, Cate 1, Cate 2, Ludwig, Curvature, and Cubic. Lee, Quad, Cate 1, and Cate 2 were used in Imbens & Kalyanaraman (2011), and Lee, Ludwig, and Curvature were used in Calonico et al. (2014); details about these datasets can be found in Imbens & Kalyanaraman 2011 (Page 18) and Calonico et al. 2014 (Page 20). We also introduce the Cubic mean function as a comparison to the Quad mean function, because in the Quad mean function the linear trends on either side of b are the opposite sign, whereas those for the Cubic mean are the same sign.

The boundary for each dataset is $b = 0$. The treatment effect is $\tau = 0.04$ for Lee and Curvature, $\tau = 0$ for Quad and Cubic, $\tau = 0.1$ for Cate 1 and 2, and $\tau = -3.35$ for Ludwig. Also displayed in Figure 1 are a set of sample points $\{(x_i, y_i), i = 1, \ldots, 500\}$ for each mean function, which shows what one dataset looks like for each mean function. One can see that—although $\epsilon_i \sim N(0, 0.1295^2)$ for all mean functions—the relative noise varies across mean functions.

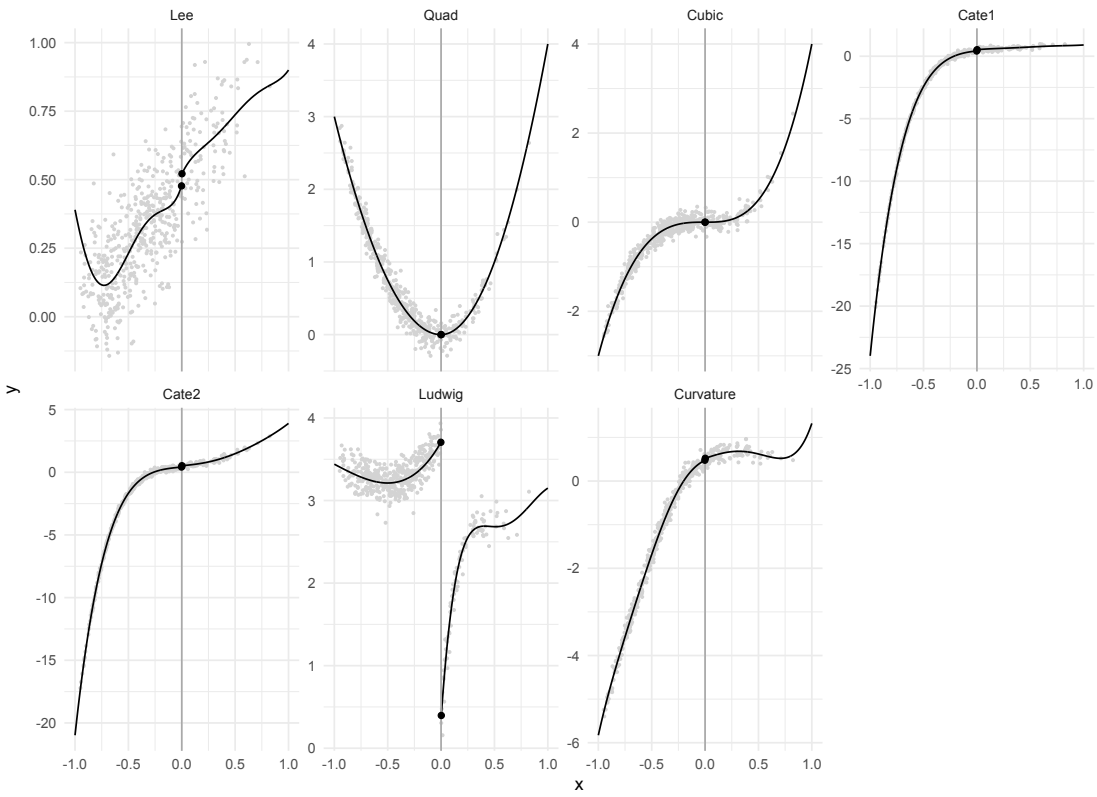

Figure 1: Mean function for the datasets used in our simulation study. Lee, Quad, Cate 1, and Cate 2 were used in Imbens & Kalyanaraman (2011), and Lee, Ludwig, and Curvature were used in Calonico et al. (2014). Displayed in gray are a set of sample points $\{(x_i, y_i), i = 1, \ldots, 500\}$ for each mean function.

For LLR and robust LLR we used the `rdrobust` `R` package (Calonico, Cattaneo, & Titiunik, 2015). For both methods, we used the bandwidth introduced in Imbens & Kalyanaraman (2011) - hereafter called the IK bandwidth - because this is the most widely used bandwidth in practice. Simulation results for other bandwidth choices for robust LLR appear in Calonico et al. (2014).

One of the motivations of our GPR methodology is that current LLR methodologies plug in a single bandwidth when there is actually uncertainty in this bandwidth choice. Thus, before we compare LLR, robust LLR, and our GPR methodology, we first ask: within this simulation study, is there uncertainty in the bandwidth choice? Figure 2 shows the distribution of the IK bandwidth used for LLR and the robust bandwidth used for the bias correction in robust LLR across the 1,000 datasets in the simulation study. Because

the underlying mean functions are the same across datasets, Figure 2 implies that there is indeed uncertainty in the bandwidth that is not being captured by either LLR or robust LLR. Furthermore, the uncertainty varies across mean functions.

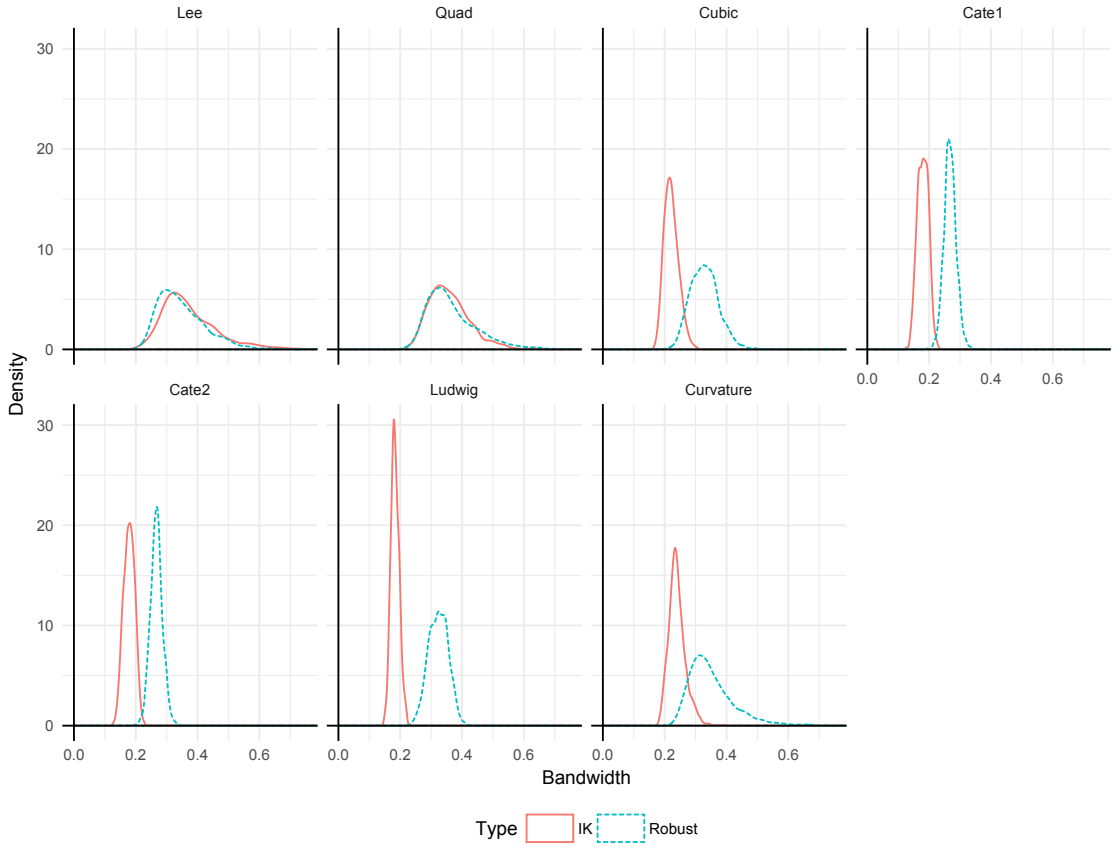

Figure 2: Selected IK bandwidths and robust bandwidths for LLR and robust LLR, respectively, across simulations. We can see that there is uncertainty in the bandwidth choice, and the amount of uncertainty varies across mean functions.

Imbens & Kalyanaraman (2011) ran a simulation study that focused on the Lee, Quad, Cate 1, and Cate 2 mean functions, and they compared different bandwidth selectors for LLR in terms of bias and root mean squared error (RMSE). Similarly, Calonico et al. (2014) ran a simulation study that focused on the Lee, Ludwig, and Curvature mean functions, and they compared different bandwidth selectors for LLR and robust LLR in terms of coverage and mean interval length (IL). We synthesize these simulation studies and report in Figure 3 how LLR, robust LLR, and two versions of our GPR methodology perform on the seven

mean functions in Figure 1 in terms of coverage, IL, absolute bias, and RMSE. The numbers plotted in Figure 3 are in Tables 1 and 2 in the Appendix.

Robust LLR is meant to improve the coverage of LLR, and indeed it does for all datasets except for the Quad dataset. Robust LLR, compared to LLR, yields confidence intervals that are shifted by a bias correction and inflated by the uncertainty in this bias correction, and so the better coverage is in part due to wider confidence intervals (see the systematically higher mean interval length at the top right of Figure 3) and in part due to better bias properties (see the bottom left of Figure 3). Robust LLR also tends to exhibit worse RMSE than LLR (see the bottom right of Figure 3).

Our primary method ("GPR") exhibits narrower confidence intervals, on average, than robust LLR. Our method also exhibits better coverage than robust LLR, except for the Lee and Ludwig datasets. Furthermore, our method exhibits lower RMSE than robust LLR - except for the Ludwig dataset - and occasionally lower RMSE than LLR. However, our method always exhibits more bias than robust LLR, which explicitly uses a bias correction.

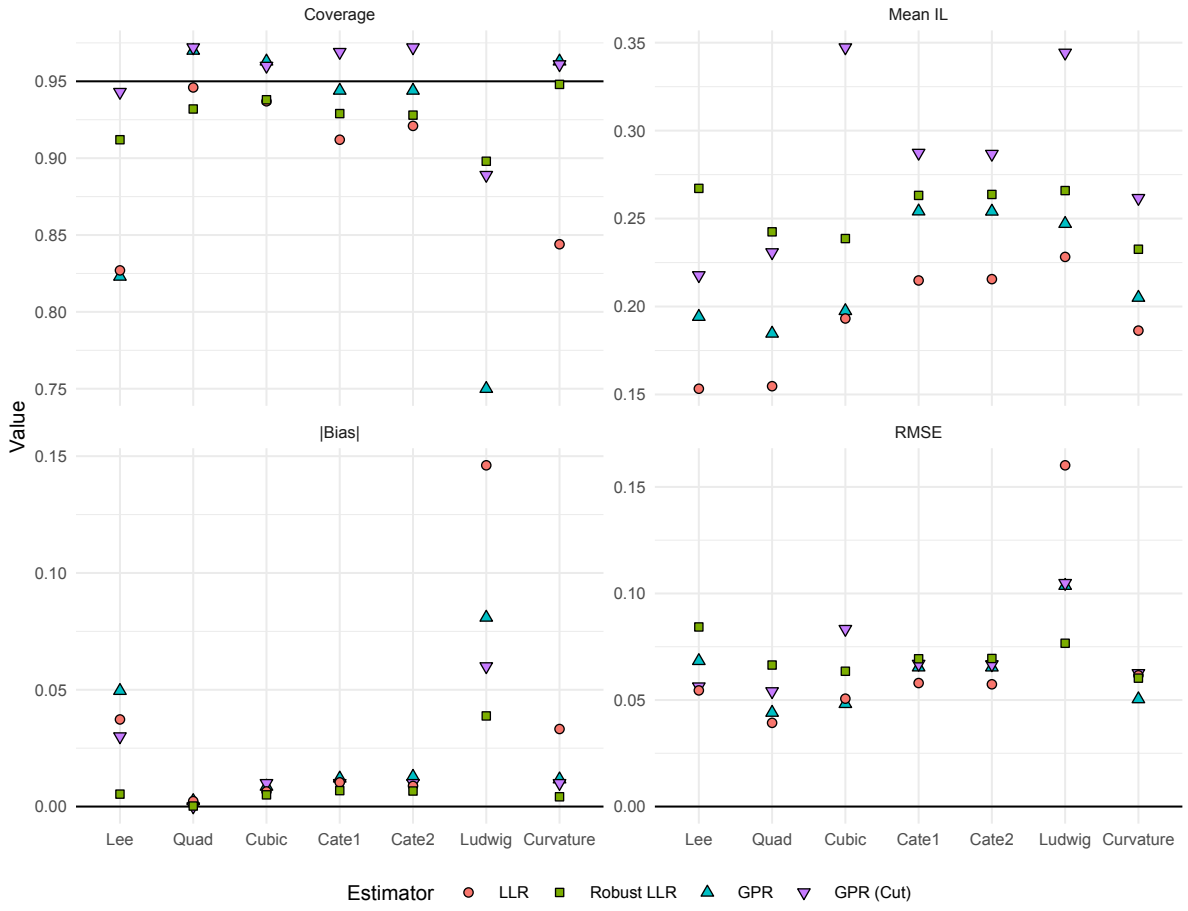

Figure 3: The coverage, mean interval length (IL), absolute bias, and root mean-squared-error (RMSE) for LLR, robust LLR, and our GPR method. The coverage for LLR on the Ludwig dataset (31.5%) is not displayed in the figure.

For the Lee and Ludwig datasets, our method did worse than robust LLR in terms of coverage and bias, and this may be because it is inappropriate to assume that $\mu_T(x)$ and $\mu_C(x)$ are stationary processes—i.e., that the covariance parameters do not vary across the running variable X—in these cases. By assuming stationarity, our GPR model uses data both close to and far from the boundary to estimate the single variance σ^2_{GP} and lengthscale ℓ. This assumption is related to the stability of the second derivative of $\mu_T(x)$ and $\mu_C(x)$, because the covariance parameters of a Gaussian process dictate their derivative processes; see Wang (2012) for a further discussion of this relationship. Figure 4 displays the absolute second derivative of the seven mean functions (in blue). The Lee and Ludwig datasets are

characterized by the absolute second derivative rapidly increasing near the boundary. Our GPR methodology likely does not do well for these datasets because we are using data far from the boundary to estimate the overall curvature of the mean function, which leads us to underestimating the curvature of the Lee and Ludwig mean functions at the boundary.

Figure 4 also displays $\frac{\hat{\sigma}_{GP}}{\hat{\ell}}$, the ratio of the maximum-likelihood estimates of the covariance parameters, which was computed within a sliding window of a noiseless version of the seven mean functions. Although there is not a one-to-one correspondence between the absolute second derivative and $\frac{\hat{\sigma}_{GP}}{\hat{\ell}}$, their behavior is notably similar, which reinforces the idea that both the variance σ^2_{GP} and lengthscale ℓ play a role in capturing the curvature of the mean function. Furthermore, this connection between the second derivative and the covariance parameters strengthens the analogy between the covariance parameters in our GPR methodology and the bandwidth in the LLR methodology, because the IK bandwidth is estimated as a nonlinear function of the estimated second derivative at the boundary b (Imbens & Kalyanaraman, 2011).

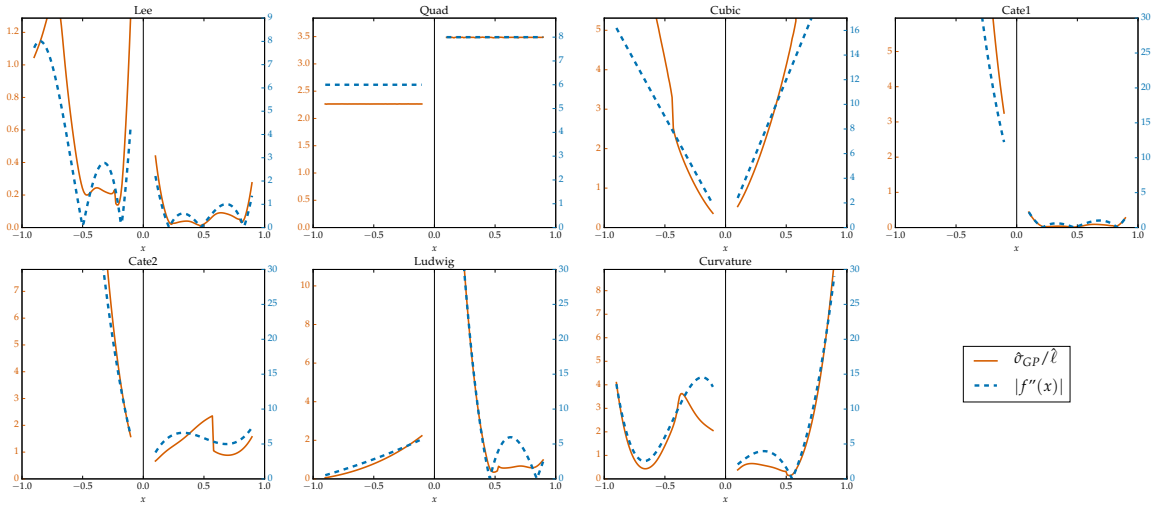

Figure 4: The absolute second derivative (blue) of the seven mean functions shown in Figure 1, and the ratio of the maximum-likelihood estimates of the variance and lengthscale (orange), computed within a sliding window of a noiseless version of the seven mean functions. The sliding window was the range $[x - 0.1, x + 1]$ for $x \in [-0.9, -0.1]$ in the control group (left-hand side) and $x \in [0.1, 0.9]$ in the treatment group (right-hand side).

If one did not believe the Stationary Assumption held, one alternative would be to only use data close to the boundary when fitting our GPR model. This is our second method, "GPR (Cut)," whose coverage, mean IL, absolute bias, and RMSE is also displayed in Figure 3. For each of the 1,000 replications of the seven mean functions, we first estimated the IK bandwidth with a rectangular kernel; then, we fit our GPR model within this estimated bandwidth. This procedure improves upon our GPR model for and only for the Lee and Ludwig datasets—the coverage increased to 94.3% and 88.9%, respectively, and the bias improved for the Lee and Ludwig datasets while staying the same for the other datasets. While the coverage also improved for the Cate 1 and Cate 2 datasets, this is likely due to the increase in the interval length. Because results improved only for these two datasets, this further demonstrates that using our GPR model on the whole dataset can be beneficial when $\mu_T(x)$ and $\mu_C(x)$ are stationary processes; otherwise, it may be preferable to only fit our GPR model to data close to the boundary.

Overall, our methodology performs well compared to LLR and robust LLR; in particular, our GPR methodology tends to yield better coverage than LLR and robust LLR when the underlying mean functions are stationary across the running variable X. The issue of undercoverage in LLR methodologies has been relatively unaddressed in the RDD literature, except for robust LLR (Calonico, Cattaneo, & Titiunik, 2014), and so our GPR methodology can be viewed as the second methodology to yield promising coverage properties for sharp RDD analyses, while also providing a flexible fit to the underlying mean functions $\mu_T(x)$ and $\mu_C(x)$.

When $\mu_T(x)$ and $\mu_C(x)$ are nonstationary across X, our GPR methodology could be extended to incorporate a lengthscale function $\ell(x)$ instead of a single lengthscale ℓ. However, such a lengthscale function would either need to be specified (see Rasmussen and Williams 2006, Chapter 4, for an example), or estimated via another Gaussian process prior (Plagemann, Kersting, & Burgard, 2008).

But this may be more than necessary: In a sharp RDD, we only need a good estimate of

the covariance parameters near the boundary, rather than across the entire mean functions $\mu_T(x)$ and $\mu_C(x)$. More work needs to be done to determine the optimal amount of data to include in our GPR model for estimating both these parameters and the LATE.

6 Discussion and Conclusion

Local linear regression (LLR) and its variants are currently the most common methodologies for estimating the local average treatment effect (LATE) in RDDs. These methods are popular because they are easy to implement and there is a large literature on their theoretical properties. However, (1) some feel limited by local linear or quadratic trends, and would like to more flexibly fit the mean treatment and control response, and (2) it is well-known that LLR tends to exhibit undercoverage in practice, but this has been relatively ignored in the RDD literature. Here, we have proposed a nonparametric Bayesian methodology for estimating the LATE in RDDs that addresses these concerns.

Rather than specifying a particular polynomial or functional form for the mean treatment and control responses, our Gaussian Process Regression (GPR) methodology flexibly fits these response functions by placing a general prior on the underlying mean treatment and control response functions. Furthermore, our methodology incorporates the uncertainty in how units' are weighted when estimating the LATE, rather than plugging in a single weighting scheme, as is the common practice with LLR methodologies.

Robust LLR (Calonico, Cattaneo, & Titiunik, 2014) was one of the first methods to address the undercoverage issue of LLR by estimating a bias correction and corresponding variance inflation that respectively shift and expand LLR's confidence intervals. We have shown via simulation that our GPR methodology tends to exhibit better coverage than both LLR and robust LLR, and also narrower confidence intervals and better RMSE properties than robust LLR. Our methodology adds to the RDD literature in general, and particularly to the literature that addresses the undercoverage issue commonly seen in RDDs.

Our work is not the first Bayesian methodology for RDDs, but - to the best of our knowledge - it is the first to propose the use of Gaussian processes for RDDs. Li et al. (2015) propose a principal stratification approach that provides alternative identification assumptions bsaed on a formal definition of local randomization. Geneletti et al. (2015) propose a Bayesian methodology that incorporates prior information in the treatment effect. Chib & Greenberg (2014) use Bayesian splines to estimate treatment effects in RDDs, which addresses the need to more flexibly fit the mean treatment and control response and exhibits better RMSE properties than LLR. Chib & Jacobi (2015) propose a Bayesian methodology specific to fuzzy RDDs. Unlike our method, none of these Bayesian methodologies specifically target the undercoverage issue of standard RDD methodologies.

Furthermore, while our methodology focuses on flexibly fitting the mean treatment and control responses while also improving coverage properties, there are many other issues of interest in the RDD literature. For example, we only consider a continuous running variable, but others have considered discrete running variables as well (e.g., Lee & Card 2008). Also, we only consider sharp RDDs; Li et al. (2015) and Chib & Jacobi (2015) provide Bayesian methodologies for fuzzy RDDs which could likely be combined with our Gaussian process approach. Furthermore, we focused on RDDs that only have one background covariate - the running variable - but other covariates could be included in our GPR methodology to improve precision of our treatment effect estimator (e.g., Calonico et al. 2016). See Imbens & Lemieux (2008) for a review of these other RDD concerns.

Finally, although our GPR methodology exhibits arguably better coverage and RMSE properties than standard methodologies in the literature, there are ways our methodology could be improved, even for the case of a sharp RDD with only one covariate. Our methodology does not perform well when the mean response functions $\mu_T(x)$ and $\mu_C(x)$ are nonstationary across the running variable X. However, we have shown that our GPR methodology performs well regardless when using only data close to the boundary. More work needs to be done to either estimate a lengthscale function $\ell(x)$ or determine the optimal amount of data

to use for our GPR methodology. Furthermore, we only focused on the squared exponential covariance function for Gaussian processes, but other covariance functions should be considered. Because estimating the LATE in an RDD is fundamentally an extrapolation issue, covariance functions whose purpose is to extrapolate well (e.g., Wilson & Adams 2013) may be particularly suitable for RDDs.

Overall, our initial Gaussian process model for RDDs exhibits promising properties compared to standard methods, and more work should be done to develop Gaussian process models that are more tailored for RDDs.

7 Appendix

7.1 Deriving the Posterior for the Shifted GP Model

Assume the Mean Shift Assumption given in Section 3.2, i.e., assume $\mu_T(x) = \mu_C(x) + \tau$. Furthermore, assume that (10) holds, i.e., the treatment response $Y_i(1)$ and control response $Y_i(0)$ have the following relationship with the running variable x_i:

$$Y_i(1) = \mu_C(x_i) + \tau + \epsilon_i, \quad \text{and } Y_i(0) = \mu_C(x_i) + \epsilon_i, \quad \text{where}$$
$$\epsilon_i \sim N(0, \sigma_y^2)$$

Finally, assume the following priors for $\mu_C(x)$ and τ:

$$\mu_C(x) \sim \text{GP}(m(\mathbf{x}), K(\mathbf{x}, \mathbf{x}'))$$
$$\tau \sim N(0, \sigma_\tau^2)$$

for fixed mean function $m(\mathbf{x})$ and covariance function $K(\mathbf{x}, \mathbf{x}')$. Then, the joint distribution of the responses $\mathbf{y}$ and the treatment effect parameter τ is

$$\begin{pmatrix} \mathbf{y} \\ \tau \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} m(\mathbf{x}) \\ 0 \end{pmatrix}, \begin{pmatrix} \text{Var}(y) & \text{Cov}(\mathbf{y}, \tau) \\ \text{Cov}(\mathbf{y}, \tau)^T & \sigma_\tau^2 \end{pmatrix}\right)$$

where

$$\begin{aligned} \text{Var}(y) &= \text{Var}(\mu_C(x) + \tau \mathbb{I}(\mathbf{x} > b) + \epsilon) \\ &= K(\mathbf{x}, \mathbf{x}') + \sigma_\tau^2 \mathbb{I}(\mathbf{x} > b)\mathbb{I}(\mathbf{x} > b)^T + \sigma_y^2 \mathbf{I}_n \\ \text{Cov}(\mathbf{y}, \tau) &= \text{Cov}(\mu_C(x) + \tau \mathbb{I}(\mathbf{x} > b) + \epsilon, \tau) \\ &= \text{Cov}(\tau \mathbb{I}(\mathbf{x} > b), \tau) \\ &= \sigma_\tau^2 \mathbb{I}(\mathbf{x} > b) \end{aligned}$$

Then, the posterior for the treatment effect as in (11) immediately follows by standard properties of conditional multivariate normal distributions.

7.2 Complete Simulation Results from Section 5

Table 1: Simulation for $n = 500$, shown in Figure 3

	Method											
	LLR				**"Robust" LLR**				**GPR**			
Dataset	**EC**	$\overline{\textbf{IL}}$	**Bias**	**RMSE**	**EC**	$\overline{\textbf{IL}}$	**Bias**	**RMSE**	**EC**	$\overline{\textbf{IL}}$	**Bias**	**RMSE**
Lee	82.7%	0.15	0.04	0.05	91.2%	0.27	0.01	0.08	82.3%	0.19	0.05	0.07
Quad	94.6%	0.15	0.00	0.04	93.2%	0.24	0.00	0.07	97.0%	0.18	0.00	0.04
Cubic	93.7%	0.19	-0.01	0.05	93.8%	0.24	0.01	0.06	96.3%	0.20	-0.01	0.05
Cate 1	91.2%	0.21	-0.01	0.06	92.9%	0.26	0.01	0.07	94.4%	0.25	-0.01	0.07
Cate 2	92.1%	0.22	-0.01	0.06	92.8%	0.26	0.01	0.07	94.4%	0.25	-0.01	0.07
Ludwig	31.5%	0.23	0.15	0.16	89.8%	0.27	0.04	0.08	75.0%	0.25	0.08	0.10
Curvature	84.4%	0.19	-0.03	0.06	94.8%	0.23	0.00	0.06	96.3%	0.21	-0.01	0.05

Simulations assessing the empirical coverage (EC), mean interval length ($\overline{\text{IL}}$), bias, and root mean-squared-error (RMSE) for local linear regression (LLR), robust LLR, and our Gaussian Process Regression (GPR) method. These methods were performed on 1,000 replications of seven different datasets, which were also used in Imbens & Kalyanaraman (2011) and Calonico et al. (2014). A plot of these numbers is shown in Figure 3.

Table 2: Simulation Results Using Only Data Within the IK Bandwidth

	Method							
	GPR				**GPR (Cut)**			
Data	**EC**	$\overline{\textbf{IL}}$	**Bias**	**RMSE**	**EC**	$\overline{\textbf{IL}}$	**Bias**	**RMSE**
Lee	82.3%	0.19	0.05	0.07	94.3%	0.22	0.03	0.06
Quad	97.0%	0.18	0.00	0.04	97.2%	0.23	-0.00	0.05
Cubic	96.3%	0.20	-0.01	0.05	96.0%	0.35	-0.01	0.08
Cate 1	94.4%	0.25	-0.01	0.07	96.9%	0.29	-0.01	0.07
Cate 2	94.4%	0.25	-0.01	0.07	97.2%	0.29	-0.01	0.07
Ludwig	75.0%	0.25	0.08	0.10	88.9%	0.34	0.06	0.10
Curvature	96.3%	0.21	-0.01	0.05	96.1%	0.26	-0.01	0.06

Simulations assessing the empirical coverage (EC), mean interval length ($\overline{\text{IL}}$), bias, and root mean-squared-error (RMSE) for GPR and GPR (Cut), which is GPR on only data inside the IK bandwidth with a rectangular kernel. Note that the GPR columns are the same as those in Table 1. GPR (Cut) performs much better than GPR for the Lee and Ludwig datasets, and GPR (Cut) performs marginally better than GPR for the Cate 1 and Cate 2 datasets. Otherwise, GPR (Cut) is equal to or worse than GPR. This demonstrates that GPR on the whole dataset can be beneficial when $\mu_T(x)$ and $\mu_C(x)$ are stationary processes; otherwise, it may be preferable to only fit our GPR model to data close to the boundary.

7.3 Assumptions for Posterior Consistency Proofs

Assumptions on the Running Variable X (Assumptions A1)

1. Assume the control running variable values $\{x_i\}_{i=1}^{n_C}$ are known elements of $[b_C, b]$ and the treatment running variable values $\{x_i\}_{i=1}^{n_T}$ are known elements of $[b, b_T]$, for some b_C, b_T, and boundary b.

Assumptions on the Response Function (Assumptions A2)

1. Let $C^\alpha[b, b_T]$ and $C^\alpha[b_C, b]$ be Hölder spaces of α-smooth functions $f : [b, b_T] \to \mathbb{R}$ and $f : [b_C, b] \to \mathbb{R}$, respectively. Assume $\mu_T(x) \in C^\alpha[b, b_T]$ and $\mu_C(x) \in C^\alpha[b_C, b]$, where $\mu_T(x)$ and $\mu_C(x)$ are the treatment and control response functions.

2. The treatment and control responses $\{y_i\}_{i=1}^{n_T}$ and $\{y_i\}_{i=1}^{n_C}$ have the following relationships with the running variable:

$$y_i = \mu_T(x_i) + \epsilon_i, \quad i = 1, \ldots, n_T \tag{20}$$

$$y_i = \mu_C(x_i) + \epsilon_i, \quad i = 1, \ldots, n_C \tag{21}$$

for mean response functions $\mu_T(x)$ and $\mu_C(x)$ and independent errors $\{\epsilon_i\}_{i=1}^{n_C} \sim N(0, \sigma_{y0}^2)$ and $\{\epsilon_i\}_{i=1}^{n_T} \sim N(0, \sigma_{y1}^2)$.

Assumptions on the Noise (Assumptions A3)

1. The priors on σ_{y0} and σ_{y1} have support on compact intervals that are subsets of $(0, \infty)$ which contain the true errors σ_{y0} and σ_{y1}, respectively.

Assumptions on the Prior for ℓ (Assumptions A4)

1. The lengthscale ℓ has a prior distribution κ such that, for positive constants C_1, D_1, C_2, D_2, nonnegative constants p, q, and every sufficiently large $a > 0$,

$$C_1 a^p \exp(-D_1 a \log^q a) \leq \kappa(a) \leq C_2 a^p \exp(-D_2 a \log^q a) \tag{22}$$

7.4 Proof of Theorems 1 and 2

Proof of Theorem 1: Assume that the Stationary Assumption holds, the covariance functions $K_T(x,x)$ and $K_C(x,x)$ are fixed, and Assumptions A1, A2, and A3 hold. Then, according to Theorem 2.2 in van der Vaart & van Zanten (2009), for sufficiently large M_T and M_C,

$$\begin{aligned} \prod(\mu_T : h(\mu_T, \mu_T^*) \geq M_T \epsilon_{n_T} | x_1, \ldots, x_{n_T}) &\xrightarrow{P_{\mu_T^*}} 0, \quad \text{and} \\ \prod(\mu_C : h(\mu_C, \mu_C^*) \geq M_C \epsilon_{n_C} | x_1, \ldots, x_{n_C}) &\xrightarrow{P_{\mu_C^*}} 0 \end{aligned} \tag{23}$$

for some contraction rates ϵ_{n_T} and ϵ_{n_C}, where h is the Hellinger distance. The nature of the contraction rates ϵ_{n_T} and ϵ_{n_C} are discussed in van der Vaart & van Zanten (2009), and depend on the differentiability and smoothness of the true $\mu_T^*(x)$ and $\mu_C^*(x)$. Note that the only covariate value for which both of these hold is $x = b$, because by Assumption A1, the intersection of the supports of $\{x_i\}_{i=1}^{n_C}$ and $\{x_i\}_{i=1}^{n_T}$ is only the boundary b.

Because the Hellinger distance is symmetric and satisfies the triangle inequality,

$$\begin{aligned} h(\tau, \tau^*) &= h\left(\mu_T(b) - \mu_C(b), \mu_T^*(b) - \mu_C^*(b)\right) \\ &\leq h(\mu_T(b), \mu_T^*(b)) + h(\mu_C(b), \mu_C^*(b)) \\ &\leq M_T \epsilon_{n_T} + M_C \epsilon_{n_C} \end{aligned}$$

where the last inequality holds with posterior probability 1, by (23). Therefore,

$$\prod(\tau : h(\tau, \tau^*) \geq M \epsilon_n | x_1, \ldots, x_n) \xrightarrow{P_{\tau^*}} 0 \tag{24}$$

where $M \equiv \sqrt{2} \cdot \max(M_T, M_C)$ and $\epsilon_n \equiv \sqrt{2} \cdot \max(\epsilon_{n_T}, \epsilon_{n_C})$, so that $M\epsilon_n \geq M_T \epsilon_{n_T} + M_C \epsilon_{n_C}$.

■

Proof of Theorem 2: Assume that the Stationary Assumption holds, the σ_{GP}^2 parameters in $K_T(x,x)$ and $K_C(x,x)$ are fixed, and Assumptions A1, A2, A3, and A4 hold.

By Theorem 3.1 of van der Vaart & van Zanten (2009), (23) holds, and then the proof of Theorem 2 is identical to that of Theorem 1. ■

7.5 Extending Theorems 1 and 2 to Other Mean and Covariance Functions and Random σ^2_{GP}

Rasmussen & Williams 2006 (Section 2.7) discuss other choices of mean functions besides $m(x) = 0$, particularly mean functions of the form $m(x) = \mathbf{h}(x)^T\boldsymbol{\beta}$ for some fixed basis functions $\mathbf{h}(x)$. However, Rasmussen & Williams (2006) argue that different choices of $m(x)$ are more for interpretability than predictive accuracy, because $m(x) = 0$ does not constrain the posterior mean to be zero, and so different choices of $m(x)$ likely do not affect the consistency results in Section 4. To the best of our knowledge, the literature has focused primarily on the choice $m(x) = 0$ for posterior consistency of GPR (e.g., Choi & Schervish 2007 and van der Vaart & van Zanten 2009).

van der Vaart & van Zanten (2009) discuss how their Theorems 2.2 and 3.1 (which correspond to our Theorems 1 and 2, respectively) extend to covariance functions besides the squared exponential. Specifically, their results hold for processes whose spectral measure has subexponential tails; the squared exponential process falls under this class, but van der Vaart & van Zanten (2009) discuss other processes that fall under this class as well, and our Theorems 1 and 2 also hold for those processes.

Finally, the literature has focused on posterior consistency for the case when a prior is placed on the lengthscale ℓ but not on the variance σ^2_{GP}, as in van der Vaart & van Zanten (2009). To the best of our knowledge, Choi (2007) is the only work to consider posterior consistency when priors are placed on both ℓ and σ^2_{GP}; however, these results only hold for binary Gaussian process regression, and their necessary assumptions are more restrictive than those presented in this paper. We leave posterior consistency of our GPR method when priors are placed on both ℓ and σ^2_{GP} as future work.

References

Calonico, S., Cattaneo, M. D., Farrell, M. H., & Titiunik, R. (2016). Regression discontinuity designs using covariates. Tech. rep., working paper, University of Michigan.

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, *82*(6), 2295–2326.

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2015). rdrobust: An r package for robust nonparametric inference in regression-discontinuity designs. *R Journal*, *7*(1), 38–51.

Card, D., Dobkin, C., & Maestas, N. (2004). The impact of nearly universal insurance coverage on health care utilization and health: evidence from medicare. Tech. rep., National Bureau of Economic Research.

Chib, S., & Greenberg, E. (2014). Nonparametric bayes analysis of the sharp and fuzzy regression discontinuity designs. Tech. rep., Mimeo 2014. http://economics. sas. upenn. edu/system/files/event_papers/cgbayesrddmarch2014. pdf (accessed February 17, 2015).

Chib, S., & Jacobi, L. (2015). Bayesian fuzzy regression discontinuity analysis and returns to compulsory schooling. *Journal of Applied Econometrics*.

Choi, T. (2007). Alternative posterior consistency results in nonparametric binary regression using gaussian process priors. *Journal of statistical planning and inference*, *137*(9), 2975–2983.

Choi, T., & Schervish, M. J. (2007). On posterior consistency in nonparametric regression problems. *Journal of Multivariate Analysis*, *98*(10), 1969–1987.

Fan, J. (1992). Design-adaptive nonparametric regression. *Journal of the American statistical Association*, *87*(420), 998–1004.

Fan, J., & Gijbels, I. (1992). Variable bandwidth and local linear regression smoothers. *The Annals of Statistics*, (pp. 2008–2036).

Gelman, A., & Imbens, G. (2014). Why high-order polynomials should not be used in regression discontinuity designs. Tech. rep., National Bureau of Economic Research.

Gelman, A., et al. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian analysis*, *1*(3), 515–534.

Geneletti, S., O'Keeffe, A. G., Sharples, L. D., Richardson, S., & Baio, G. (2015). Bayesian regression discontinuity designs: Incorporating clinical knowledge in the causal analysis of primary care data. *Statistics in medicine*, *34*(15), 2334–2352.

Ghosal, S., & Roy, A. (2006). Posterior consistency of gaussian process prior for nonparametric binary regression. *The Annals of Statistics*, (pp. 2413–2429).

Hahn, J., Todd, P., & Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, *69*(1), 201–209.

Imbens, G., & Kalyanaraman, K. (2011). Optimal bandwidth choice for the regression discontinuity estimator. *The Review of economic studies*, (p. rdr043).

Imbens, G. W., & Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of econometrics*, *142*(2), 615–635.

Lee, D. S., & Card, D. (2008). Regression discontinuity inference with specification error. *Journal of Econometrics*, *142*(2), 655–674.

Li, F., Mattei, A., Mealli, F., et al. (2015). Evaluating the causal effect of university grants on student dropout: evidence from a regression discontinuity design using principal stratification. *The Annals of Applied Statistics*, *9*(4), 1906–1931.

Ludwig, J., & Miller, D. L. (2005). Does head start improve children's life chances? evidence from a regression discontinuity design. Tech. rep., National Bureau of Economic Research.

Matsudaira, J. D. (2008). Mandatory summer school and student achievement. *Journal of Econometrics*, *142*(2), 829–850.

Plagemann, C., Kersting, K., & Burgard, W. (2008). Nonstationary gaussian process regression using point estimates of local smoothness. *Machine learning and knowledge discovery in databases*, (pp. 204–219).

Porter, J. (2003). Estimation in the regression discontinuity model. *Unpublished Manuscript, Department of Economics, University of Wisconsin at Madison*, (pp. 5–19).

Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press.

Rubin, D. B. (1977). Assignment to treatment group on the basis of a covariate. *Journal of educational Statistics*, *2*(1), 1–26.

Thistlethwaite, D. L., & Campbell, D. T. (1960). Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational psychology*, *51*(6), 309–317.

Van der Klaauw, W. (2002). Estimating the effect of financial aid offers on college enrollment: A regression–discontinuity approach. *International Economic Review*, *43*(4), 1249–1287.

van der Vaart, A. W., & van Zanten, J. H. (2008). Rates of contraction of posterior distributions based on gaussian process priors. *The Annals of Statistics*, (pp. 1435–1463).

van der Vaart, A. W., & van Zanten, J. H. (2009). Adaptive bayesian estimation using a gaussian random field with inverse gamma bandwidth. *The Annals of Statistics*, (pp. 2655–2675).

Wang, X. (2012). *Bayesian modeling using latent structures*. Ph.D. thesis, Citeseer.

Wilson, A. G., & Adams, R. P. (2013). Gaussian process kernels for pattern discovery and extrapolation. In *ICML (3)*, (pp. 1067–1075).